Optimal Patching in Clustered Malware Epidemics

S. Eshghi, MHR. Khouzani, S. Sarkar, S. S. Venkatesh

arXiv:1403.1639v1 [cs.CR] 7 Mar 2014

Abstract

Studies on the propagation of malware in mobile networks have revealed that the spread of malware can be highly inhomogeneous. Platform diversity, contact list utilization by the malware, clustering in the network structure, etc. can also lead to differing spreading rates. In this paper, a general formal framework is proposed for leveraging such heterogeneity to derive optimal patching policies that attain the minimum aggregate cost due to the spread of malware and the surcharge of patching. Using Pontryagin's Maximum Principle for a stratified epidemic model, it is analytically proven that in the mean-field deterministic regime, optimal patch disseminations are simple single-threshold policies. Through numerical simulations, the behavior of optimal patching policies is investigated in sample topologies and their advantages are demonstrated.

Index Terms

Security, Wireless Networks, Immunization and Healing, Belief Propagation, Technology Adoption

I. Introduction

Worms (self-propagating malicious codes) are a decades-old threat in the realm of the Internet. Worms undermine the network in various ways: they can eavesdrop on and analyze traversing data, access privileged information, hijack sessions, disrupt network functions such as routing, etc. Although the Internet is the traditional arena for trojans, spyware, and viruses, the current boom in mobile devices, combined with their spectacular software and hardware capabilities, has created a tremendous opportunity for future malware. Mobile devices communicate with each other and with computers through myriad means– Bluetooth or Infrared when they are in close proximity, multimedia messages (MMS), mobile Internet, and peer to peer networks. Current smartphones are equipped with operating systems, CPUs, and memory powerful enough to execute complex codes. Wireless malware such as *cabir*, *skulls*, *mosquito*, *commwarrior* have already sounded the alarm [1]. It has been theoretically predicted [2] that it is only a matter of time before major malware outbreaks are witnessed in the wireless domain.

Malware spreads when an infective node *contacts*, i.e., communicates with, a susceptible node, i.e., a node without a copy of the malware and vulnerable to it. This spread can be countered through patching [3]: the vulnerability utilized by the worm can be fixed by installing security patches that immunize the susceptible and potentially remove the malware from the infected, hence simultaneously healing and immunizing infective nodes. However, the distribution of these patches burdens the limited resources of the network, and can wreak havoc on the system if not carefully controlled. In wired networks, the spread of *Welchia*, a counter-worm to thwart *Blaster*, rapidly destabilized important sections of the Internet [4]. Resource constraints are even more pronounced in wireless networks, where bandwidth is more sensitive to overload and nodes have limited energy reserves. Recognizing the above, works such as [5], [6] have included the cost of patching in the aggregate damage of the malware and have characterized the optimal dynamic patching policies that attain desired trade-offs between the patching efficacy and the extra taxation of network resources. These results, however, critically rely on the homogeneous mixing assumption: that all pairs of nodes have identical expected inter-contact times. Optimality may now be attained by policies that constrain the nodes in the same state to take the same action. While this assumption may serve as an approximation in cases where detailed information about the network is unavailable, studies [2], [7]–[10] show that the spread of malware in mobile networks can be very inhomogeneous, owing primarily to the non-uniform distribution of nodes. Thus, a uniform action may be sub-optimal.

We can motivate heterogeneous epidemics in various ways: 1) Proximity-based spread, where locality causes heterogeneity; 2) Software/protocol diversity [11]–[13], which has even been envisioned as a defense mechanism against future threats [11]; 3) IP space diversity, e.g., worms that use the IP masks of specific Autonomous Systems (AS) to increase their yield [14]; 4) Differing clique sizes, especially for epidemics for which there is an underlying social network graph [3], [15]; 5) Behavioral patterns, specifically how much risky behaviour every agent engages in [16]; 6) Clusters in cloud computing [17]; 7) Technology adoption, belief-formation over social media, and health care [18]–[22].

Indeed, many works have proposed practical methods to identify, characterize and incorporate such inhomogeneities to more accurately predict the spread of infection [2], [8], [9], [14], [23]–[25], etc. Relatively few, e.g., [3], [10], [12], consider the cost of patching and seek to minimize it in the presence of heterogeneous contact processes. The proposed policies in [3], [10] are heuristic and apply to specific settings. The only paper we could find that provides *provably optimal* patching policies

S. Eshghi, S. Sarkar and S. S. Venkatesh are with the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA, U.S.A. Their email addresses are *eshghi,swati,venkates@seas.upenn.edu*. MHR. Khouzani is with the Department of Electrical Engineering at the University of Southern California, Los Angeles, CA. His e-mail address is *rezaeikh@usc.edu*.

This paper was presented [in part] at the IEEE Information Theory and Applications Workshop (ITA '12) , San Diego, CA, February, 2012

for heterogeneous networks is [12]. They, however, focus on SIS models and optimize only in the space of static policies (those that do not vary patching rates over time) therein. Patching performance can be significantly improved if we allow the patching rates to vary dynamically in accordance with the evolution of the infection. Characterization of the optimal controls in the space of dynamic and clustered policies has, however, so far remained elusive.

We propose a formal framework for deriving *dynamic optimal* patching policies that leverage heterogeneity in the network structure to attain the minimum possible aggregate cost due to the spread of malware and the overhead of patching. We assume arbitrary (potentially non-linear) functions for the cost rates of the infective nodes. We consider both *non-replicative* and *replicative* patching: in the former, some of the hosts are pre-loaded with the patch, which they transmit to the rest. In the latter, each recipient of the patch can also forward the patch to nodes that it contacts by a mechanism similar to the spread of the malware itself. In our model, patching can immunize susceptible nodes and may or may not heal infective nodes. The framework in each case relies on optimal control formulations that cogently capture the effect of the patching rate controls on the state dynamics and their resulting trade-offs. We accomplish this by using a combination of damage functions associated with the controls and a *stratified*[1] mean-field deterministic epidemic model in which nodes are divided into different types. Nodes of the same type homogeneously mix with a rate specific to that type, and nodes of different types contact each other at rates particular to that pair of types. The model can therefore capture any communication topology between different groups of nodes. Above and beyond, it can exploit the inhomogeneity in the network to enable a better utilization of the resources. Such higher patching efficacy is achieved by allowing the patching controls to depend on node types, which in turn leads to multidimensional (dynamic) optimal control formulations. We first develop our system dynamics and objectives (§II) and characterize optimal non-replicative (§III) and replicative (§IV) patching. We then analyze an alternate objective (§V) and present numerical simulation of our results (§VI).

Multidimensional optimal control formulations, particularly those in the solution space of functions rather than variables, are usually associated with the pitfall of amplifying the complexity of the optimization. An important contribution of the paper, therefore, is to prove that for both non-replicative and replicative settings the optimal control associated with each type has a simple structure provided the corresponding patching cost is either concave or convex. Furthermore, our analysis, using Pontryagin's Maximum Principle, reveals that the structure of the optimal control for a specific type depends only on the nature of the corresponding patching cost and not on those of other types. This holds even though the control for each type affects immunization and healing in other types and the spread of the infection in general. Specifically, if the patching cost associated with the control for a given type is concave, irrespective of the nature of the patching costs for other types, the corresponding optimal control turns out to be a bang-bang function with at most one jump: up to a certain threshold time (possibly different for different types) it selects the maximum possible patching rate and subsequently it stops patching altogether. If the patching cost is strictly convex, the decrease from the maximum to the minimum patching rate is continuous rather than abrupt, and monotonous. To the best of our knowledge, such simple structure results have not been established in the context of (static or dynamic) control of heterogeneous epidemics. Our numerical calculations reveal a series of interesting behaviors of optimal patching policies for different sample topologies.

II. System Model and Objective Formulation

In this section we describe and develop the model of the state dynamics of the system as a general *stratified* epidemic for both non-replicative (§II-A) and replicative (§II-B) patching, motivate the model (§II-C), formulate the aggregate cost of patching, and cast this resource-aware patching as a multi-dimensional optimal control problem (§II-E). This formulation relies on a key property of the state dynamics which we isolate in §II-D. We develop solutions in this model framework and present our main results in §III and §IV.

Our models are based on mean-field limits of Poisson contact processes for which pathwise convergence results have been shown (c.f. [26, p.1], [27]).

A. Dynamics of non-replicative patching

A node is **infective** if it has been contaminated by the malware, **susceptible** if it is vulnerable to the infection but not yet infected, and **recovered** if it is immune to the malware. An infective node spreads the malware to a susceptible one while transmitting data or control messages. The network consists of nodes that can be stratified into M different *types* (equivalently, *clusters, segments, populations, categories, classes, strata*). The population of these types need not be equal. Nodes of type i contacts those of type j at rate β_{ij}.

For type i, $S_i(t)$, $I_i(t)$, and $R_i(t)$ are respectively the fraction of susceptible, infective and recovered states at time t. Therefore, for all t and all i, we have $S_i(t) + I_i(t) + R_i(t) = 1$. We assume that during the course of the epidemic, the population of each type is stable and does not change with time.

Amongst each type, a pre-determined set of nodes, called *dispatchers*, are preloaded with the appropriate patch. Dispatchers can transmit patches to both susceptible and infective nodes, *immunizing* the susceptible and possibly *healing* the infective; in

[1]Known by other terms such as structured, clustered, multi-class, multi-type, multi-population, compartmental epidemic models, and sometimes loosely as heterogeneous, inhomogeneous or spatial epidemic models.

either case successful transmission converts the target node to the recovered state. In *non-replicative* patching (as opposed to *replicative* patching- see §II-B) the recipient nodes of the patch do not propagate it any further.[2] Dispatchers of type i contact nodes of type j at rate $\bar{\beta}_{ij}$, which may be different from the malware contact rate β_{ij} between these two types. Examples where contact rates may be different include settings where the network manager may utilize a higher priority option for the distribution of patches, ones where the malware utilizes legally restricted means of propagation not available to dispatchers, or ones where the patch is not applicable to all types, with the relevant $\bar{\beta}_{ij}$ now being zero. The fraction of dispatchers in type i, which is fixed over time in the non-replicative setting, is R_i^0, where $0 \leq R_i^0 < 1$.

Place the time origin $t = 0$ at the earliest moment the infection is detected and the appropriate patches generated. Suppose that at $t = 0$, for each i, an initial fraction $0 \leq I_i(0) = I_i^0 \leq 1$ of nodes of type i are infected; we set $I_i^0 = 0$ if the infection does not initially exist amongst a type i. At the onset of the infection, the dispatchers are the only agents immune to the malware, hence constituting the initial population of recovered nodes. We therefore identify $R_i(0) = R_i^0$. In view of node conservation, it follows that $S_i^0 = 1 - I_i^0 - R_i^0$ represents the initial fraction $S_i(0)$ of susceptible nodes of type i.

At any given t, susceptibles of type i may be contacted by infectives of type j at rate β_{ji}. We may fold resistance to infection into the contact rates and so from a modeling perspective, we may assume that susceptibles contacted by infectious agents are instantaneously infected. Accordingly, susceptibles of type i are transformed to infectives (of the same type) at an aggregate rate of $S_i(t) \sum_{j=1}^{M} \beta_{ji} I_j(t)$ by contact with infectives of any type.

The system manager regulates the resources consumed in the patch distribution by dynamically controlling the rate at which dispatchers contact susceptible and infective nodes. For each j, let the control function $u_j(t)$ represent the rate of transmission attempts of dispatchers of type j at time t. We suppose that the controls are non-negative and bounded,

$$0 \leq u_j(\cdot) \leq u_{j,\max}. \tag{1}$$

We will restrict consideration to control functions $u_j(\cdot)$ that have a finite number of points of discontinuity. We say that a control (vector) $\mathbf{u}(t) = \big(u_1(t), \ldots, u_M(t)\big)$ is *admissible* if each $u_j(t)$ has a finite number of points of discontinuity.

Given the control $\mathbf{u}(t)$, susceptibles of type i are transformed to recovered nodes of the same type at an aggregate rate of $S_i(t) \sum_{j=1}^{M} \bar{\beta}_{ji} R_j^0 u_j(t)$ by contact with dispatchers of any type. A subtlety in the setting is that the dispatcher may find that the efficacy of the patch is lower when treating infective nodes. This may model situations, for instance, where the malware attempts to prevent the reception or installation of the patch in an infective host, or the patch is designed only to remove the vulnerability that leaves nodes exposed to the malware but does not remove the malware itself if the node is already infected. We capture such possibilities by introducing a (type-dependent) coefficient $0 \leq \pi_{ji} \leq 1$ which represents the efficacy of patching an infective node: $\pi_{ji} = 0$ represents one extreme where a dispatcher of type j can only immunize susceptibles but can not heal infectives of type i, while $\pi_{ji} = 1$ represents the other extreme where contact with a dispatcher of type j both immunizes and heals nodes of type i equally well; we also allow π_{ij} to assume intermediate values between the above extremes. An infective node transforms to the recovered state if a patch heals it; otherwise, it remains an infective. Infective nodes of type i accordingly recover at an aggregate rate of $I_i(t) \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j^0 u_j(t)$ by contact with dispatchers.

We say that a type j is a *neighbour* of a type i if $\beta_{ij} > 0$, and $S_j > 0$ (i.e., infected nodes of type i can contact nodes of type j). There is now a natural notion of a topology that is inherited from these rates with types as vertices and edges between neighboring types. Figure 1 illustrates some simple topologies. For a given topology inherited from the rates

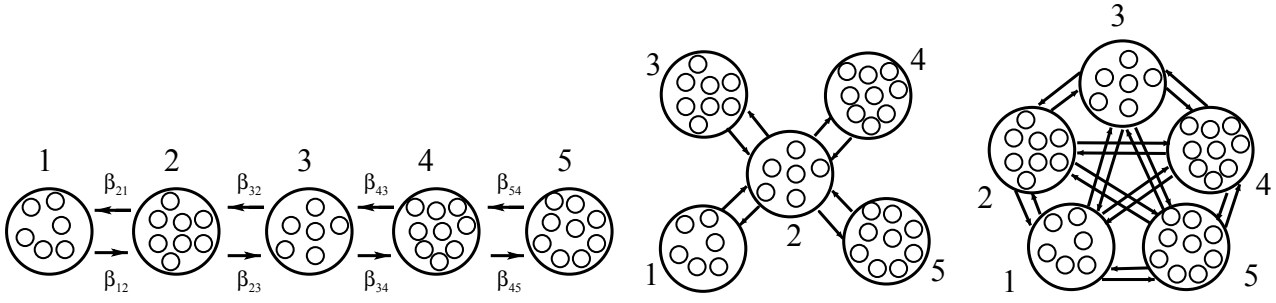

Fig. 1. Three sample topologies of 5 hotspot regions: linear, star and complete. For instance, nodes of hotspot 1 in the linear topology can only communicate with nodes of hotspots 1 and 2: they contact nodes of hotspot 1 at rate β_{11} and nodes of hotspot 2 at rate β_{12}.

$\{\beta_{ij}, 1 \leq i, j \leq M\}$ there is now another natural notion, that of *connectivity*: we say that type j is connected to type i if, for some k, there exists a sequence of types $i = s_1 \mapsto s_2 \mapsto \ldots \mapsto s_{k-1} \mapsto s_k = j$ where type s_{l+1} is a neighbour of type s_l for $1 \leq l < k$. We assume that each type is either initially infected ($I_i(0) > 0$), or is connected to an initially infected type. We also assume that for every type i such that $R_i^0 > 0$, there exists a type j for which $\bar{\beta}_{ij} > 0$, i.e., type i can immunize nodes

[2]This may be preferred if the patches themselves can be contaminated and cannot be reliably authenticated.

of at least one type, and there exist types k and l for which $\beta_{ki} > 0$ and $\beta_{il} > 0$, $S_l(0) > 0$, i.e., the infection can spread to and from that type. (In most settings we may expect, naturally, that $\beta_{ii} > 0$ and $\bar{\beta}_{ii} > 0$.)

Thus, we have:[3]

$$\dot{S}_i = -\sum_{j=1}^{M} \beta_{ji} I_j S_i - S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j^0 u_j, \tag{2a}$$

$$\dot{I}_i = \sum_{j=1}^{M} \beta_{ji} I_j S_i - I_i \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j^0 u_j, \tag{2b}$$

where, by writing $\mathbf{S}(t) = \big(S_1(t), \ldots, S_M(t)\big)$, $\mathbf{I}(t) = \big(I_1(t), \ldots, I_M(t)\big)$, and $\mathbf{R}^0(t) = \big(R_1^0(t), \ldots, R_M^0(t)\big)$ in a compact vector notation, the initial conditions and state constraints are given by

$$\mathbf{S}(0) = \mathbf{S}^0 \succeq \mathbf{0}, \quad \mathbf{I}(0) = \mathbf{I}^0 \succeq \mathbf{0}, \tag{3}$$

$$\mathbf{S}(t) \succeq \mathbf{0}, \quad \mathbf{I}(t) \succeq \mathbf{0}, \quad \mathbf{S}(t) + \mathbf{I}(t) \preceq \mathbf{1} - \mathbf{R}^0. \tag{4}$$

In these expressions $\mathbf{0}$ and $\mathbf{1}$ represent vectors all of whose components are 0 and 1 respectively, and the vector inequalities are to be interpreted as component-wise inequalities. Note that the evolution of $\mathbf{R}(t)$ need not be explicitly considered since at any given time, node conservation gives $R_i(t) = 1 - S_i(t) - I_i(t)$. We henceforth drop the dependence on t and make it implicit whenever we can do so without ambiguity.

B. Dynamics of replicative patching

In the replicative setting, a recipient of the patch can forward it to other nodes upon subsequent contact. Thus, recovered nodes of type i are added to the pool of dispatchers of type i, whence the fraction of dispatchers of type i grows from the initial $R_i(0) = R_i^0$ to $R_i(t)$ at time t. This should be contrasted with the non-replicative model in which the fraction of dispatchers of type i is fixed at R_i^0 for all t.

The system dynamics equations given in (2) for the non-replicative setting now need to be modified to take into account the growing pool of dispatchers. While in the non-replicative case we chose the pair $\big(\mathbf{S}(t), \mathbf{I}(t)\big)$ to represent the system state, in the replicative case it is slightly more convenient to represent the system state by the explicit triple $\big(\mathbf{S}(t), \mathbf{I}(t), \mathbf{R}(t)\big)$. The system dynamics are now governed by:

$$\dot{S}_i = -\sum_{j=1}^{M} \beta_{ji} I_j S_i - S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j u_j, \tag{5a}$$

$$\dot{I}_i = \sum_{j=1}^{M} \beta_{ji} I_j S_i - I_i \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j u_j, \tag{5b}$$

$$\dot{R}_i = S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j u_j + I_i \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j u_j, \tag{5c}$$

with initial conditions and state constraints given by

$$\mathbf{S}(0) = \mathbf{S}^0 \succeq \mathbf{0}, \quad \mathbf{I}(0) = \mathbf{I}^0 \succeq \mathbf{0}, \quad \mathbf{R}(0) = \mathbf{R}^0 \succeq \mathbf{0}, \tag{6}$$

$$\mathbf{S}(t) \succeq \mathbf{0}, \quad \mathbf{I}(t) \succeq \mathbf{0}, \quad \mathbf{R}(t) \succeq \mathbf{0}, \quad \mathbf{S}(t) + \mathbf{I}(t) + \mathbf{R}(t) = \mathbf{1}. \tag{7}$$

The assumptions on controls and connectivity are as in §II-A.

C. Motivation of the models and instantiations

We now motivate the stratified epidemic models (2) and (5) through different examples, instantiating the different types in each context.

1) Proximity-based spread—heterogeneity through locality: The overall roaming area of the nodes can be divided into regions (e.g., hotspots, office/residential areas, central/peripheral areas, etc.) of different densities (fig. 1). One can therefore stratify the nodes based on their locality, i.e., each type corresponds to a region. IP eavesdropping techniques (using software such as `AirJack`, `Ethereal`, `FakeAP`, `Kismet`, etc.) allow malware to detect new victims in the vicinity of the host. Distant nodes have more attenuated signal strength (i.e., lower SINR) and are therefore less likely to be detected. Accordingly, malware (and also patch) propagation rates β_{ij} (respectively $\bar{\beta}_{ij}$) are related to the local densities of the nodes in each region and decay with an increase in the distance between regions i and j: typically β_{ii} exceeds β_{ij} for $i \neq j$, likewise for $\bar{\beta}_{ij}$. The same phenomenon was observed for malware such as `cabir` and `lasco` that use Bluetooth and Infrared to propagate.

[3]We use dots to denote *time* derivatives throughout, e.g., $\dot{S}(t) = dS(t)/dt$.

2) Heterogeneity through software/protocol diversity: A network that relies on a homogeneous software/protocol is vulnerable to an attack that exploits a common weakness (e.g., a buffer overflow vulnerability). Accordingly, inspired by the natural observation that *the chances of survival are improved by heterogeneity*, increasing the network's heterogeneity without sacrificing interoperability has been proposed as a defense mechanism [11]. In practice, mobile nodes use different operating systems and communication protocols, e.g., `Symbian`, `Android`, `IOS`, `RIM`, `webOS`, etc. Such heterogeneities lead to dissimilar rates of propagation of malware amongst different types, where each type represents a specific OS, platform, software, protocol, etc. In the extreme case, the malware may not be able to contaminate nodes of certain types. The patching response should take such inhomogeneities into account in order to optimally utilize network resources, since the rate of patching can also be dissimilar among different types.

3) Heterogeneity through available IP space: Smartphone trojans like `skulls` and `mosquito` spread using Internet or P2P networks. In such cases the network can be decomposed into *autonomous systems* (ASs) with each type representing an AS [14]. A worm either scans IP addresses uniformly randomly or uses the IP masks of ASs to restrict its search domain and increase its rate of finding new susceptible nodes. In each of these cases the contact rates differ between different AS's depending on the actual number of assigned IPs in each IP sub-domain and the maximum size of that IP sub-domain.

4) Heterogeneity through differing clique sizes: Malware that specifically spreads in social networks has been recorded in the past few years [15]. Examples include `Samy` in MySpace in 2005 and `Koobface` in MySpace and Facebook in 2008. `Koobface`, for instance, spread by delivering (contaminated) messages to the "friends" of an infective user. MMS based malware such as `commwarrior` can also utilize the contact list of an infective host to access new handsets. In such cases the social network graph can be approximated by a collection of friendship *cliques*.[4] Users of the same clique can be regarded as the same type with the rate of contact within cliques and across cliques differing depending on the relative sizes of the cliques.

5) Cloud-computing—heterogeneity through cluster sizes: In cluster (or grid, or volunteer) computing [17], each cluster of CPUs in the cloud constitutes a type. Any two computers in the same cluster can communicate at faster rates than those in different clusters. These contact rates depend on the communication capacity of connecting lines as well as the relative number of computers in each cluster.

6) Clustered epidemics in technology adoption, belief-formation over social media and health care: We now elaborate on the application of our clustered epidemics model in these diverse set of contexts. First consider a rivalry between two technologies or companies for adoption in a given population, e.g., Android and iPhone, or cable and satellite television. Individuals who are yet to choose either may be considered as susceptibles and those who have chosen one or the other technology would be classified as either infective or recovered depending upon their choice. Dispatchers constitute the promoters of a given technology (the one whose subscribers are denoted as recovered). Awareness about the technology and subsequent subscription to either may spread through social contact between infectives and susceptibles (infection propagation in our terminology), and dispatchers and the rest (patching in our terminology). Immunization of a susceptible corresponds to her adoption of the corresponding technology, while healing of an infective corresponds to an alteration in her original choice. The stratifications may be based on location or social cliques, and the control $\mathbf{u}$ would represent promotion efforts, which would be judiciously selected by the proponent of the corresponding technology. Patching may either be replicative or non-replicative depending on whether the newly subscribed users are enticed to attract more subscribers by referral rewards. Similarly, clustered epidemics may be used to model belief management over social media, where infective and recovered nodes represent individuals who have conflicting persuasions and susceptibles represent those who are yet to subscribe to either doctrine. Last, but not least, the susceptible-infective-recovered classification and immunization/healing/infection have natural connotations in the context of a biological epidemic. Here, the dispatchers correspond to health-workers who administer vaccines and/or hospitalization and the stratification is based on location. Note that in this context, patching can only be non-replicative.

D. Key observations

A natural but important observation is that if the initial conditions are non-negative, then the system dynamics (2) and (5) yield unique states satisfying the positivity and normalisation constraints (4) and (7), respectively. The proof is technical and is not needed elsewhere in the paper; we relegate it accordingly to the appendix.

Theorem 1. *The dynamical system* (2) *(respectively* (5)*) with initial conditions* (3) *(respectively,* (6)*) has a unique state solution* $\big(\mathbf{S}(t), \mathbf{I}(t)\big)$ *(respectively* $\big(\mathbf{S}(t), \mathbf{I}(t), \mathbf{R}(t)\big)$*) which satisfies the state constraints* (4) *(respectively,* (7)*). For all* $t > 0$*,* $I_i(t) > 0$*;* $R_i(t) > 0$ *if* $R_i(0) > 0$*; and* $S_j(t) > 0$ *if and only if* $S_j(0) \neq 0$*. For each* j *such that* $S_j(0) = 0$*,* $S_j(t') = 0$ *for all* $t' \in (0, T]$*.*

E. The optimality objective

The network seeks to minimize the overall cost of infection and the resource overhead of patching in a given operation time window $[0, T]$. At any given time t, the system incurs costs at a rate $f\big(\mathbf{I}(t)\big)$ due to the malicious activities of the malware. For instance, the malware may use infected hosts to eavesdrop, analyze, misroute, alter, or destroy the traffic that the hosts

[4] A clique is a maximal complete subgraph of a graph [28, p. 112].

generate or relay. We suppose that $f(\mathbf{0}) = 0$ and make the natural assumption that the scalar function $f(\mathbf{I})$ is increasing and differentiable with respect to each I_i. The simplest natural candidate for $f(\mathbf{I})$ is of the form $\sum_{i=1}^{M} f_i(I_i)$; in this setting each f_i is a non-decreasing function of its argument representing the cost of infection for type i which is in turn determined by the criticality of the data of that type and its function.[5] The network also benefits at the rate of $L\big(\mathbf{R}(t)\big)$, i.e., incurs a cost at the rate of $-L\big(\mathbf{R}(t)\big)$, due to the removal of uncertainty about the state of the nodes being patched. We suppose that the scalar function $L(\mathbf{R})$ is non-decreasing and differentiable with respect to each R_i.

In addition to the cost of infection, each dispatcher burdens the network with a cost by consuming either available bandwidth, energy reserves of the nodes (e.g., in communication and computing networks), or money (e.g., in technology adoption, propaganda, health-care) to disseminate the patches. Suppose dispatchers of type i incur cost at a rate of $R_i^0 h_i(u_i)$. We suppose that the overhead of extra resource (bandwidth or energy or money) consumption at time t is then given by a sum of the form $\sum_{i=1}^{M} R_i^0 h_i(u_i)$. The scalar functions $h_i(\cdot)$ represent how much resource is consumed for transmission of the patch by nodes of each type and how significant this extra taxation of resources is for each type. Naturally enough, we assume these functions are non-decreasing and satisfy $h_i(0) = 0$ and $h_i(\gamma) > 0$ for $\gamma > 0$. We assume, additionally, that each h_i is twice differentiable. Following the same lines of reasoning, the corresponding expression for the cost of replicative patching is of the form $\sum_{i=1}^{M} R_i h_i(u_i)$.

With the arguments stated above, the aggregate cost for non-replicative patching is given by an expression of the form

$$J_{non-rep} = \int_0^T \left(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_i^0 h_i(u_i) \right) dt, \tag{8}$$

while for replicative patching, the aggregate cost is of the form

$$J_{rep} = \int_0^T \left(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_i h_i(u_i) \right) dt. \tag{9}$$

Problem Statement: The system seeks to minimize the aggregate cost **(A)** in (8) for non-replicative patching (2, 3) and **(B)** in (9) for replicative patching (5, 6) by an appropriate selection of an optimal admissible control $\mathbf{u}(t)$.

In this setting, it is clear that by a scaling of $\bar{\beta}_{ji}$ and $h_j(\cdot)$ we can simplify our admissibility conditions on the optimal controls by supposing that $u_{j,\max} = 1$.

It is worth noting that any control in the non-replicative case can be emulated in the replicative setting: this is because the fraction of the dispatchers in the replicative setting is non-decreasing, hence at any time instance, a feasible $u^{rep}(t)$ can be selected such that $R_i(t)u_i^{rep}(t)$ is equal to $R_i^0 u_i^{non-rep}(t)$. This means that the minimum cost of replicative patching is always less than the minimum cost of its non-replicative counterpart. Our numerical results will show that this improvement is substantial. However, replicative patches increase the risk of patch contamination: the security of a smaller set of known dispatchers is easier to manage than that of a growing set whose identities may be ambiguous. Hence, in a nutshell, if there is a dependable mechanism for authenticating patches, replicative patching is the preferred method, otherwise one needs to evaluate the trade-off between the risk of compromised patches and the efficiency of the patching.

III. Optimal Non-Replicative Patching

A. Numerical framework for computing the optimal controls

The main challenge in computing the optimal state and control functions $\big((\mathbf{S}, \mathbf{I}), \mathbf{u}\big)$ is that while the differential equations (2) can be solved once the optimal controls $\mathbf{u}(\cdot)$ are known, an exhaustive search for an optimal control is infeasible as there are an uncountably infinite number of control functions. *Pontryagin's Maximum Principle (PMP)* provides an elegant technique for solving this seemingly intractable problem (c.f. [29]). Referring to the integrand in (8) as $\xi_{non-rep}$ and the RHS of (2a) and (2b) as ν_i and μ_i, we define the Hamiltonian to be

$$\mathcal{H} = \mathcal{H}(\mathbf{u}) := \xi_{non-rep} + \sum_{i=1}^{M} (\lambda_i^S \nu_i + \lambda_i^I \mu_i), \tag{10}$$

where the *adjoint* (or *costate*) functions λ_i^S and λ_i^I are continuous functions that for each $i = 1 \dots M$, and at each point of continuity of $\mathbf{u}(\cdot)$, satisfy

$$\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S_i}, \quad \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I_i}, \tag{11}$$

along with the final (i.e., transversality) conditions

$$\lambda_i^S(T) = 0, \quad \lambda_i^I(T) = 0. \tag{12}$$

[5] Such differences themselves may be a source of stratification. In general, different types need not exclusively reflect disparate mixing rates.

Then PMP implies that the optimal control at time t satisfies

$$\mathbf{u} \in \arg\min_{\mathbf{v}} \mathcal{H}(\mathbf{v}) \tag{13}$$

where the minimization is over the space of admissible controls (i.e., $H(u) = \min_v H(v)$).

In economic terms, the adjoint functions represent a shadow price (or imputed value); they measure the marginal worth of an increment in the state at time t when moving along an optimal trajectory. Intuitively, in these terms, λ_i^I ought to be positive as it represents the additional cost that the system incurs per unit time with an increase in the fraction of infective nodes. Furthermore, as an increase in the fraction of the infective nodes has worse long-term implications for the system than an increase in the fraction of the susceptibles, we anticipate that $\lambda_i^I - \lambda_i^S > 0$. The following result confirms this intuition. It is of value in its own right but as its utility for our purposes is in the proof of our main theorem of the following section, we will defer its proof (to §III-C) to avoid breaking up the flow of the narrative at this point.

Lemma 1. *The positivity constraints $\lambda_i^I(t) > 0$ and $\lambda_i^I(t) - \lambda_i^S(t) > 0$ hold for all $i = 1, \ldots, M$ and all $t \in [0, T)$.*

The abstract maximum principle takes on a very simple form in our context. Using the expression for $\xi_{non-rep}$ from (8) and the expressions for ν_i and μ_i from (2), trite manipulations show that the minimization (13) may be expressed in the much simpler (nested) scalar formulation

$$u_i(t) \in \arg\min_{0 \le x \le 1} \psi_i(x, t) \qquad (1 \le i \le M); \tag{14}$$

$$\psi_i(x, t) := R_i^0(h_i(x) - \phi_i(t)x); \tag{15}$$

$$\phi_i := \sum_{j=1}^{M} \bar{\beta}_{ij}\lambda_j^S S_j + \sum_{j=1}^{M} \bar{\beta}_{ij}\pi_{ij}\lambda_j^I I_j. \tag{16}$$

Equation (14) allows us to characterise u_i as a function of the state and adjoint functions at each time instant. Plugging into (2) and (11), we obtain a system of (non-linear) differential equations that involves only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), and where the initial values of the states (3) and the final values of the adjoint functions (12) are known. Numerical methods for solving *boundary value* nonlinear differential equation problems may now be used to solve for the state and adjoint functions corresponding to the optimal control, thus providing the optimal controls using (14).

We conclude this section by proving an important property of $\phi_i(\cdot)$, which we will use in subsequent sections.

Lemma 2. *For each i, $\phi_i(t)$ is a decreasing function of t.*

Proof: We examine the derivative of $\phi_i(t)$; we need expressions for the derivatives of the adjoint functions towards that end. From (10), (11), at any t at which $\mathbf{u}$ is continuous, we have:

$$\begin{aligned}
\dot{\lambda}_i^S &= -\frac{\partial L(\mathbf{R})}{\partial R_i} - (\lambda_i^I - \lambda_i^S)\sum_{j=1}^{M} \beta_{ji} I_j + \lambda_i^S \sum_{j=1}^{M} \bar{\beta}_{ji} R_j^0 u_j, \\
\dot{\lambda}_i^I &= -\sum_{j=1}^{M} \Big((\lambda_j^I - \lambda_j^S)\beta_{ij} S_j \Big) + \lambda_i^I \sum_{j=1}^{M} \pi_{ji}\bar{\beta}_{ji} R_j^0 u_j - \frac{\partial L(\mathbf{R})}{\partial R_i} - \frac{\partial f(\mathbf{I})}{\partial I_i}.
\end{aligned} \tag{17}$$

Using (16), (17) and some reassembly of terms, at any t at which $\mathbf{u}$ is continuous,

$$\dot{\phi}_i(t) \;=\; -\sum_{j=1}^{M} \bar{\beta}_{ij} \Bigg[S_j \frac{\partial L(\mathbf{R})}{\partial R_j} \;+\; \pi_{ij} I_j \left(\frac{\partial L(\mathbf{R})}{\partial R_j} + \frac{\partial f(\mathbf{I})}{\partial I_j} \right) \;+\; \sum_{k=1}^{M} (1 \;+\; \pi_{ij}) \lambda_j^I \beta_{kj} S_j I_k \;+\; \sum_{k=1}^{M} \pi_{ij} I_j S_k \beta_{jk} (\lambda_k^I \;-\; \lambda_k^S) \Bigg].$$

The assumptions on $L(\cdot)$ and $f(\cdot)$ (together with Theorem 1) show that the first two terms inside the square brackets on the right are always non-negative. Theorem 1 and Lemma 1 (together with our assumptions on π_{ij}, β_{ij} and $\bar{\beta}_{ij}$) show that the penultimate term is positive for $t > 0$ and the final term is non-negative. It follows that $\dot{\phi}_i(t) < 0$ for every $t \in (0, T)$ at which $\mathbf{u}(t)$ is continuous. As $\phi_i(t)$ is a continuous function of time and its derivative is negative except at a finite number of points (where $\mathbf{u}$ may be discontinuous), it follows indeed that, as advertised, $\phi_i(t)$ is a decreasing function of time. ■

B. Structure of optimal non-replicative patching

We are now ready to identify the structures of the optimal controls $(u_1(t), \ldots, u_M(t))$:

Theorem 2. *Predicated on the existence of an optimal control, for types i such that $R_i^0 > 0$: if $h_i(\cdot)$ is concave, then the optimal control for type i has the following structure: $u_i(t) = 1$ for $0 < t < t_i$, and $u_i(t) = 0$ $t_i < t \le T$, where $t_i \in [0, T)$. If $h_i(\cdot)$ is strictly convex then the optimal control for type i, $u_i(t)$ is continuous and has the following structure: $u_i(t) = 1$ for $0 < t < t_i^1$, $u_i(t) = 0$ for $t_i^2 < t \le T$, and $u_i(t)$ strictly decreases in the interval $[t_i^1, t_i^2]$, where $0 \le t_i^1 < t_i^2 \le T$.*

Notice that if $R_i^0 = 0$ in (2), the control u_i is irrelevant and can take any arbitrary admissible value. Intuitively, at the onset of the epidemic, a large fraction of nodes are susceptible to the malware ("potential victims"). Bandwidth and power resources should hence be used maximally in the beginning (in all types), rendering as many infective and susceptible nodes robust against the malware as possible. In particular, there is no gain in deferring patching since the efficacy of healing infective nodes is less than that of immunizing susceptible nodes (recall that $\pi_{ij} \leq 1$). While the non-increasing nature of the optimal control is intuitive, what is less apparent is the characteristics of the decrease, which we establish in this theorem. For concave $h_i(\cdot)$, nodes are patched at the maximum possible rate until a time instant when patching stops abruptly, while for strictly convex $h_i(\cdot)$, this decrease is continuous. It is instructive to note that the structure of the optimal action taken by a type only depends on its own patching cost and not on that of its neighbours. This is somewhat counter-intuitive as the controls for one type affect the infection and recovery of other types. The timing of the decrease in each type differs and depends on the location of the initial infection as well as the topology of the network, communication rates, etc.

Proof: For non-linear concave $h_i(\cdot)$, (14) requires the minimization of the (non-linear concave) difference between a non-linear concave function of a scalar variable x and a linear function of x at all time instants; hence the minimum can only occur at the end-points of the interval over which x can vary. Thus all that needs to be done is to compare the values of $\psi_i(x,t)$ for the following two candidates: $x = 0$ and $x = 1$. Note that $\psi_i(0,t) = 0$ at all time instants and $\psi_i(1,t)$ is a function of time t. Let

$$\gamma_i(t) := \psi_i(1,t) = R_i^0 h_i(1) - R_i^0 \phi_i(t). \tag{18}$$

Then the optimal u_i satisfies the following condition:

$$u_i(t) = \begin{cases} 1 & \gamma_i(t) < 0 \\ 0 & \gamma_i(t) > 0 \end{cases} \tag{19}$$

From the transversality conditions in (12) and the definition of $\phi_i(t)$ in (16), for all i, it follows that $\phi_i(T) = 0$. From the definition of the cost term, $h_i(1) > 0$, hence, since $R_i^0 > 0$, therefore $\gamma_i(T) > 0$. Thus the structure of the optimal control predicted in the theorem for the strictly concave case will follow from (19) if we can show that $\gamma_i(t)$ is an increasing function of time t, as that implies that it can be zero at most at one point t_i, with $\gamma_i(t) < 0$ for $t < t_i$ and $\gamma_i(t) > 0$ for $t > t_i$. From (18), γ_i will be an increasing function of time if ϕ_i is a decreasing function of time, a property which we showed in Lemma 2.

If $h_i(\cdot)$ is linear (i.e., $h_i(x) = K_i x$, $K_i > 0$, since $h_i(x) > 0$ for $x > 0$), $\psi_i(x,t) = R_i^0 x(K_i - \phi_i(t))$ and from (14), the condition for an optimal u_i is:

$$u_i(t) = \begin{cases} 1 & \phi_i(t) > K_i \\ 0 & \phi_i(t) < K_i \end{cases} \tag{20}$$

But from (12), $\phi_i(T) = 0 < K_i$ and as by Lemma 2, $\phi_i(t)$ is decreasing, it follows that $\phi_i(t)$ will be equal to K_i at most at one time instant $t = t_i$, with $\phi_i(t) > 0$ for $t < t_i$ and $\phi_i(t) < 0$ for $t > t_i$. This, along with (20), concludes the proof of the theorem for the concave case.

We now consider the case where $h_i(\cdot)$ is strictly convex. In this case, the minimization in (14) may also be attained at an interior point of $[0,1]$ (besides 0 and 1) at which the partial derivative of the right hand side with respect to x is zero. Hence,

$$u_i(t) = \begin{cases} 1 & 1 < \eta(t) \\ \eta(t) & 0 < \eta(t) \leq 1 \\ 0 & \eta(t) \leq 0. \end{cases} \tag{21}$$

where $\eta(t)$ is such that $\left.\dfrac{dh_i(x)}{dx}\right|_{(x=\eta(t))} = \phi_i(t)$.

Note that $\phi_i(t)$ is a continuous function due to the continuity of the states and adjoint functions. We showed that it is also a decreasing function of time (Lemma 2). Since $h_i(\cdot)$ is double differentiable, its first derivative is continuous, and since it is strictly convex, its derivative is a strictly increasing function of its argument. Therefore, $\eta(t)$ must be a continuous and decreasing function of time, as per the predicted structure. ■

C. Proof of Lemma 1

Proof: From (17) and (12), at time T we have:

$$\lambda_i^I|_{t=T} = (\lambda_i^I - \lambda_i^S)|_{t=T} = 0,$$
$$\lim_{t\uparrow T} \dot{\lambda}_i^I = -\frac{\partial L(\mathbf{R})}{\partial R_i}(T) - \frac{\partial f(\mathbf{I})}{\partial I_i}(T) < 0$$
$$\lim_{t\uparrow T} (\dot{\lambda}_i^I - \dot{\lambda}_i^S) = -\frac{\partial f(\mathbf{I})}{\partial I_i}(T) < 0$$

Hence, $\exists\ \epsilon > 0$ s.t. $\lambda_i^I > 0$ and $(\lambda_i^I - \lambda_i^S) > 0$ over $(T-\epsilon, T)$.

Now suppose that, going backward in time from $t = T$, (at least) one of the inequalities is first violated at $t = t^*$ for i^*, i.e., for all i, $\lambda_i^I(t) > 0$ and $(\lambda_i^I(t) - \lambda_i^S(t)) > 0$ for all $t > t^*$ and either (A) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^S(t^*)) = 0$ or (B) $\lambda_{i^*}^I(t^*) = 0$ for some $i = i^*$. Note that from continuity of the adjoint functions $\lambda_i^I(t^*) \geq 0$ and $(\lambda_i^I(t^*) - \lambda_i^S(t^*)) \geq 0$ for all i.

We investigate case (A) first. We have:[6,7]

$$(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) = -\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j] - \lambda_{i^*}^I \sum_{j=1}^{M} \bar{\beta}_{ji^*}(1 - \pi_{ji^*})R_j^0 u_j.$$

First of all, $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$. The other two terms are non-positive, due to the definition of t^* and $\pi_{ij} \leq 1$. Hence, $(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) < 0$, which is in contradiction with Property 1 of real-valued functions, proved in [30]:

Property 1. *Let $g(t)$ be a continuous and piecewise differentiable function of t. If $g(t_0) = L$ and $g(t) > L$ ($g(t) < L$) for all $t \in (t_0, t_1]$. Then $\dot{g}(t_0^+) \geq 0$ (respectively $\dot{g}(t_0^+) \leq 0$).*

On the other hand, for case (B) we have:[7]

$$\dot{\lambda}_{i^*}^I(t^{*+}) = -\frac{\partial L(\mathbf{R})}{\partial R_{i^*}} - \frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j],$$

which is negative since $-\partial L(\mathbf{R})/\partial R_{i^*} \leq 0$, $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$, and from the definition of t^* (for the third term). This contradicts Property 1 and the claim follows. ∎

IV. Optimal Replicative Patching

A. Numerical framework for computing the optimal controls

As in the non-replicative setting, we develop a numerical framework for calculation of the optimal solutions using PMP, and then we establish the structure of the optimal controls.

For every control $\tilde{\mathbf{u}}$, we define $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}}) \in [0, T]$ as follows: If $R_i(0) > 0$, and therefore $R_i(t) > 0$ for all $t > 0$ due to Theorem 1, we define $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}})$ to be 0. Else, $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}})$ is the maximum t for which $R_i(t) = 0$. It follows from Theorem 1 that $R_i(t) = 0$ for all $t \leq \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}})$ and all i such that $R_i(0) = 0$, and $R_i(t) > 0$ for all $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}}) < t \leq T$. We begin with the hypothesis that there exists at least one optimal control, say $\tilde{\mathbf{u}} \in \mathcal{U}^*$, and construct a control $\mathbf{u}$ that chooses $u_i(t) := 0$ for $t \leq \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}})$ and $u_i(t) := \tilde{u}_i(t)$ for $t > \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}})$. Clearly, the states $\mathbf{S}(t), \mathbf{I}(t), \mathbf{R}(t)$ corresponding to $\tilde{\mathbf{u}}$ also constitute the state functions for $\mathbf{u}$, as the state equations only differ at $t = 0$, a set of measure zero. Thus, $\mathbf{u}$ is also an optimal control, and $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}}) = \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \mathbf{u})$ for each i. Henceforth, for notational convenience, we will refer to $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \tilde{\mathbf{u}}), \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \mathbf{u})$ as τ_i. Note that the definition of this control completely specifies the values of each u_i in $[0, T]$.

Referring to the integrand of (9) as ξ_{rep} and the RHS of equations (5a,b,c) as ν_i, μ_i and ρ_i the Hamiltonian becomes:

$$\mathcal{H} = \mathcal{H}(\mathbf{u}) := \xi_{rep} + \sum_{i=1}^{M}[(\lambda_i^S \nu_i + \lambda_i^I \mu_i + \lambda_i^R \rho_i), \tag{22}$$

where the *adjoint* functions $\lambda_i^S, \lambda_i^I, \lambda_i^R$ are continuous functions that at each point of continuity of $\mathbf{u}(\cdot)$ and for all $i = 1 \ldots M$, satisfy

$$\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S_i},\ \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I_i},\ \dot{\lambda}_i^R = -\frac{\partial \mathcal{H}}{\partial R_i}, \tag{23}$$

with the final constraints:

$$\lambda_i^S(T) = \lambda_i^I(T) = \lambda_i^R(T) = 0. \tag{24}$$

[6] $g(t_0^+) := \lim_{t\downarrow t_0} g(t)$ and $g(t_0^-) := \lim_{t\uparrow t_0} g(t)$.

[7] The RHS of the equation is evaluated at $t = t^*$ due to continuity.

According to PMP, any optimal controller must satisfy:

$$\mathbf{u} \in \arg\min_{\mathbf{v}} \mathcal{H}(\mathbf{v}), \tag{25}$$

where the minimization is over the set of admissible controls.

Using the expressions for ξ_{rep} from (9) and the expressions for ν_i, μ_i and ρ_i from (5), it can be shown that the vector minimization (25) can be expressed as a scalar minimization

$$u_i(t) \in \arg\min_{0\le x\le 1} \psi_i(x,t) \qquad (1 \le i \le M); \tag{26}$$

$$\psi_i(x,t) := R_i(t)(h_i(x) - \phi_i(t)x); \tag{27}$$

$$\phi_i := \sum_{j=1}^{M} \bar{\beta}_{ij}(\lambda_j^S - \lambda_j^R)S_j + \sum_{j=1}^{M} \pi_{ij}\bar{\beta}_{ij}(\lambda_j^I - \lambda_j^R)I_j. \tag{28}$$

Equation (26) characterizes the optimal control u_i as a function of the state and adjoint functions at each instant. Plugging the optimal u_i into the state and adjoint function equations (respectively (5) and (23)) will again leave us with a system of (non-linear) differential equations that involves only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), the initial values of the states (6) and the final values of the adjoint functions (24). Similar to the non-replicative case, the optimal controls may now be obtained (via (26)) by solving the above system of differential equations.

We conclude this subsection by stating and proving some important properties of the adjoint functions (Lemma 3 below) and $\phi_i(\cdot)$ (Lemma 4 subsequently), which we use later.

First, from (24), $\psi_i(0,t) = 0$, hence (26) results in $\psi_i(u_i,t) \le 0$. Furthermore, from the definition of τ_i, if $t \le \tau_i$, $(h_i(u_i(t)) - \phi_i(t)u_i(t)) = 0$, and if $t > \tau_i$, $(h_i(u_i(t)) - \phi_i(t)u_i(t)) = \dfrac{\psi_i(u_i,t)}{R_i(t)} \le 0$, so for all t,

$$\alpha_i(u_i,t) := (h_i(u_i(t)) - \phi_i(t)u_i(t)) \le 0. \tag{29}$$

Lemma 3. *For all $t \in [0,T)$ and for all i, we have $(\lambda_i^I - \lambda_i^S) > 0$ and $(\lambda_i^I - \lambda_i^R) > 0$.*

Using our previous intuitive analogy, Lemma 3 implies that infective nodes are always worse for the evolution of the system than either susceptible or healed nodes, and thus the marginal price of infectives is greater than that of susceptible and healed nodes at all times before T. As before, we defer the proof of this lemma (to §IV-C) to avoid breaking up the flow of the narrative. We now state and prove Lemma 4.

Lemma 4. *For each i, $\phi_i(t)$ is a decreasing function of t, and $\dot{\phi}_i(t^+) < 0$ and $\dot{\phi}_i(t^-) < 0$ for all t.*

Proof: $\phi_i(t)$ is continuous everywhere (due to the continuity of the states and adjoint functions) and differentiable whenever $\mathbf{u}(\cdot)$ is continuous. At any t at which $\mathbf{u}(\cdot)$ is continuous, we have:

$$\dot{\phi}_i(t) = \sum_{j=1}^{M} \bar{\beta}_{ij}[(\dot{\lambda}_j^S - \dot{\lambda}_j^R)S_j + (\lambda_j^S - \lambda_j^R)\dot{S}_j + \pi_{ij}(\dot{\lambda}_j^I - \dot{\lambda}_j^R)I_j + \pi_{ij}(\lambda_j^I - \lambda_j^R)\dot{I}_j].$$

From (22) and the adjoint equations (23), at points of continuity of the control, we have:

$$\begin{aligned}
\dot{\lambda}_i^S =& -(\lambda_i^I - \lambda_i^S)\sum_{j=1}^{M} \beta_{ji}I_j - (\lambda_i^R - \lambda_i^S)\sum_{j=1}^{M} \bar{\beta}_{ji}R_j u_j,\\
\dot{\lambda}_i^I =& -\frac{\partial f(\mathbf{I})}{\partial I_i} - \sum_{j=1}^{M}(\lambda_j^I - \lambda_j^S)\beta_{ij}S_j - (\lambda_i^R - \lambda_i^I)\sum_{j=1}^{M} \pi_{ji}\bar{\beta}_{ji}R_j u_j,\\
\dot{\lambda}_i^R =& \frac{\partial L(\mathbf{R})}{\partial R_i} + u_i\sum_{j=1}^{M} \bar{\beta}_{ij}(\lambda_j^S - \lambda_j^R)S_j + u_i\sum_{j=1}^{M} \pi_{ij}\bar{\beta}_{ij}(\lambda_j^I - \lambda_j^R)I_j - h_i(u_i) = \frac{\partial L(\mathbf{R})}{\partial R_i} - \alpha_i(u_i,t).
\end{aligned} \tag{30}$$

Therefore, after some regrouping and cancellation of terms, at any t, we have

$$-\dot{\phi}_i(t^+) = \sum_{j=1}^{M} \bar{\beta}_{ij}[(1-\pi_{ij})\sum_{k=1}^{M}(\lambda_j^I - \lambda_j^R)\beta_{kj}I_k S_j + \pi_{ij}\frac{\partial f(\mathbf{I})}{\partial I_j}I_j + (S_j + \pi_{ij}I_j)(\frac{\partial L(\mathbf{R})}{\partial R_j} - \alpha_i(u_i,t)) + \pi_{ij}I_j\sum_{k=1}^{M}(\lambda_k^I - \lambda_k^S)\beta_{jk}S_k].$$

Now, since $0 \le \pi_{ij} \le 1$, the assumptions on $\bar{\beta}_{ij}$, β_{ki} and β_{il}, Theorem 1, and Lemma 3 all together imply that the sum of the first and last terms of the RHS will be positive. The second and third terms will be non-negative due to the definitions of $f(\cdot)$ and $L(\cdot)$ and (29). So $\dot{\phi}_i(t^+) < 0$ for all t. The proof for $\dot{\phi}_i(t^-) < 0$ is exactly as above. In a very similar fashion, it can be proved that $\dot{\phi}_i(t) < 0$ at all points of continuity of $\mathbf{u}(\cdot)$, which coupled with the continuity of $\phi_i(t)$ shows that it is a decreasing function of time. ■

B. Structure of optimal replicative dispatch

Theorem 3. *If an optimal control exists, for types i such that $R_i(t) > 0$ for some t: if $h_i(\cdot)$ is concave for type i, the optimal control for type i has the following structure: $u_i(t) = 1$ for $0 < t < t_i$, and $u_i(t) = 0$ for $t_i < t \leq T$, where $t_i \in [0, T)$. If $h_i(\cdot)$ is strictly convex, the optimal control for type i, $u_i(t)$ is continuous and has the following structure: $u_i(t) = 1$ for $0 < t < t_i^1$, $u_i(t) = 0$ for $t_i^2 < t \leq T$, and $u_i(t)$ strictly decreases in the interval $[t_i^1, t_i^2]$, where $0 \leq t_i^1 < t_i^2 \leq T$.*

Notice that for i such that $R_i(t) = 0$ for all t, the control $u_i(t)$ is irrelevant and can take any arbitrary value. We first prove the theorem for $t \in [\tau_i, T]$, and then we show that $\tau_i \in \{0, T\}$, completing our proof.

Proof: First consider an i such that $h_i(\cdot)$ is concave and non-linear. Note that hence $\psi_i(x, t)$ is a non-linear concave function of x. Thus, the minimum can only occur at extremal values of x, i.e., $x = 0$ and $x = 1$. Now $\psi_i(0, t) = 0$ at all times t, so to obtain the structure of the control, we need to examine $\psi_i(1, t)$ at each $t > \tau_i$. Let $\gamma_i(t) := \psi_i(1, t) = R_i(t)(h_i(1) - \phi_i(t))$ be a function of time t. From (26), the optimal u_i satisfies:

$$u_i(t) = \begin{cases} 1 & \gamma_i(t) < 0, \\ 0 & \gamma_i(t) > 0. \end{cases} \tag{31}$$

We now show that $\gamma_i(t) > 0$ for an interval $(t_i, T]$ for some t_i, and $\gamma_i(t) < 0$ for $[\tau_i, t_i)$ if $t_i > \tau_i$. From (24) and (28), $\gamma_i(T) = h_i(1)R_i(T) > 0$. Since $\gamma_i(t)$ is a continuous function of its variable (due to the continuity of the states and adjoint functions), it will be positive for a non-zero interval leading up to $t = T$. If $\gamma_i(t) > 0$ for all $t \in [\tau_i, T]$, the theorem follows. Otherwise, from continuity, there must exist a $t = t_i > \tau_i$ such that $\gamma_i(t_i) = 0$. We show that for $t > t_i$, $\gamma_i(t) > 0$, from which it follows that $\gamma_i(t) < 0$ for $t < t_i$ (by a contradiction argument). The theorem will then follow from (31).

Towards establishing the above, we show that $\dot{\gamma}_i(t^+) > 0$ and $\dot{\gamma}_i(t^-) > 0$ for any t such that $\gamma_i(t) = 0$. Hence, there will exist an interval $(t_i, t_i + \epsilon)$ over which $\gamma_i(t) > 0$. If $t_i + \epsilon \geq T$, then the claim holds, otherwise there exists a $t = t_i^{'} > t_i$ such that $\gamma_i(t_i^{'}) = 0$ and $\gamma_i(t) \neq 0$ for $t_i < t < t_i^{'}$ (from the continuity of $\gamma_i(t)$). So $\dot{\gamma}_i(t_i^{'-}) > 0$, which contradicts a property of real-valued functions (proved in [30]), establishing the claim:

Property 2. *If $g(x)$ is a continuous and piecewise differentiable function over $[a, b]$ such that $g(a) = g(b)$ while $g(x) \neq g(a)$ for all x in (a, b), $\frac{dg}{dx}(a^+)$ and $\frac{dg}{dx}(b^-)$ cannot be positive simultaneously.*

We now show that $\dot{\gamma}_i(t^+) > 0$ and $\dot{\gamma}_i(t^-) > 0$ for any $t > \tau_i$ such that $\gamma_i(t) = 0$. Due to the continuity of $\gamma_i(t)$ and the states, and the finite number of points of discontinuity of the controls, for any $t > \tau_i$ we have:

$$\dot{\gamma}_i(t^+) = (\dot{R}_i(t^+)\frac{\gamma_i(t)}{R_i(t)} - R_i(t)\dot{\phi}_i(t^+)) \tag{32}$$

$$\dot{\gamma}_i(t^-) = (\dot{R}_i(t^-)\frac{\gamma_i(t)}{R_i(t)} - R_i(t)\dot{\phi}_i(t^-)). \tag{33}$$

If $\gamma_i(t) = 0$, then $\dot{\gamma}_i(t^+) = -R_i(t)\dot{\phi}_i(t^+)$ and $\dot{\gamma}_i(t^-) = -R_i(t)\dot{\phi}_i(t^-)$, which are both positive from Lemma 4 and Theorem 1, and thus the theorem follows.

The proofs for linear and strictly convex $h_i(\cdot)$'s are virtually identical to the corresponding parts of the proof of Theorem 2 and are omitted for brevity; the only difference is that in the linear case we need to replace R_i^0 with $R_i(t)$. The following lemma, proved in §IV-D, completes the proof of the theorem.

Lemma 5. *For all $0 \leq i \leq B$, $\tau_i \in \{0, T\}$.* ■

C. Proof of Lemma 3

Proof: First, from (28) and (24), we have $\phi_i(T) = 0$, which, combined with (27) results in $\psi_i(x, T) = R_i(T)h_i(x)$. Since either $R_i(T) > 0$ or $\tau_i > T$, (26) and the definition of u_i result in $u_i(T) = 0$, as all other values of x would produce a positive $\psi_i(x, T)$. Therefore, $h_i(u_i(T)) = 0$.

The rest of the proof has a similar structure to that of Lemma 1. $(\lambda_i^I - \lambda_i^S)|_{t=T} = 0$ and $\lim_{t\uparrow T}(\dot{\lambda}_i^I - \dot{\lambda}_i^S) = -\partial f(\mathbf{I})/\partial I_i < 0$, for all i. Also, for all i, $(\lambda_i^I - \lambda_i^R)|_{t=T} = 0$ and $\lim_{t\uparrow T}(\dot{\lambda}_i^I - \dot{\lambda}_i^R) = -\partial f(\mathbf{I})/\partial I_i - \partial L(\mathbf{R})/\partial R_i + h_i(u_i(T)) < 0$, since $h_i(u_i(T)) = 0$.

Hence, $\exists\ \epsilon > 0$ such that $(\lambda_i^I - \lambda_i^S) > 0$ and $(\lambda_i^I - \lambda_i^R) > 0$ over $(T - \epsilon', T)$.

Now suppose that (at least) one of the inequalities is first[8] violated at $t = t^*$ for i^*, i.e., for all i, $(\lambda_i^I(t) - \lambda_i^S(t)) > 0$ and $(\lambda_i^I(t) - \lambda_i^R(t)) > 0$ for all $t > t^*$, and either (A) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^S(t^*)) = 0$, or (B) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^R(t^*)) = 0$ for some i^*. Note that from continuity of the adjoint functions, $(\lambda_i^I(t^*) - \lambda_i^S(t^*)) \geq 0$, and $(\lambda_i^I(t^*) - \lambda_i^R(t^*)) \geq 0$ for all i.

[8]Going backward in time from $t = T$.

Case (A): Here, we have:[9]

$$(\dot{\lambda}^I_{i^*} - \dot{\lambda}^S_{i^*})(t^{*+}) = -\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}(\lambda^I_j - \lambda^S_j)\beta_{i^*j}S_j - (\lambda^I_{i^*} - \lambda^R_{i^*})\sum_{j=1}^{M}\bar{\beta}_{ji^*}(1-\pi_{ji^*})R_j u_j.$$

First of all, $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$. Also, the second and third terms are non-positive, according to the definition of t^*. Hence, $(\dot{\lambda}^I_{i^*} - \dot{\lambda}^S_{i^*})(t^{*+}) < 0$, which contradicts Property 1, therefore case (A) does not arise.

Case (B): In this case, we have:[9]

$$(\dot{\lambda}^I_{i^*} - \dot{\lambda}^R_{i^*})(t^{*+}) = -\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \frac{\partial L(\mathbf{R})}{\partial R_{i^*}} - \sum_{j=1}^{M}(\lambda^I_j - \lambda^S_j)\beta_{i^*j}S_j + \alpha_{i^*}(u_{i^*}, t).$$

We have $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$ and $-\partial L(\mathbf{R})/\partial R_{i^*} \leq 0$. The term $-(\lambda^I_{i^*} - \lambda^S_{i^*})\sum_{j=1}^{M}\beta_{ji^*}S_j$ is non-positive, according to the definition of t^*, and α_{i^*} will be non-negative due to (29). This shows $(\dot{\lambda}^I_{i^*} - \dot{\lambda}^R_{i^*})(t^{*+}) < 0$, contradicting Property 1, and so case (B) does not arise either, completing the proof. ■

D. Proof of Lemma 5

Proof: We start by creating another control $\bar{\mathbf{u}}$ from $\mathbf{u}$ such that for every i, for every $t \leq \tau_i$, $\bar{u}_i(t) := 1$, and for every $t > \tau_i$, $\bar{u}_i(t) := u_i(t)$. We prove by contradiction that $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \bar{\mathbf{u}}) \in \{0, T\}$ for each i. Since $\bar{u}_i \not\equiv u_i$ only in $[0, \tau_i]$ and $R_i(t) = 0$ for $t \in (0, \tau_i]$ when $\mathbf{u}$ is used, the state equations can only differ at a solitary point $t = 0$, and therefore both controls result in the same state evolutions. Thus, for each i, $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \bar{\mathbf{u}}) = \tau_i(\mathbf{I}(0), \mathbf{S}(0), \mathbf{R}(0), \mathbf{u})$, and $\tau_i(\mathbf{I}(0), \mathbf{S}(0), \bar{\mathbf{u}})$ may be denoted as τ_i as well. The lemma therefore follows.

For the contradiction argument, assume that the control is $\bar{\mathbf{u}}$ and that $\tau_i \in (0, T)$ for some i. Our proof relies on the fact that if $\bar{u}_i(t') = 0$ at some $t' \in (0, T)$, then $\bar{u}_i(t) = 0$ for $t > t'$, which follows from the definition of $\bar{\mathbf{u}}$ and prior results in the proof of Theorem 3.

For $t \in [0, \tau_i]$, since $R_i(t) = 0$ in this interval, (5c) becomes $\dot{R}_i = \sum_{j=1, j\neq i}^{M}\bar{\beta}_{ji}(S_i + \pi_{ji}I_i)R_j\bar{u}_j = 0$ in this interval. Since all terms in $\sum_{j=1, j\neq i}^{M}\bar{\beta}_{ji}(S_i + \pi_{ji}I_i)R_j\bar{u}_j$ are non-negative, for each $j \neq i$ we must either have (i) $\bar{\beta}_{ji}(S_i(t) + \pi_{ji}I_i(t)) = 0$ for some $t \in [0, \tau_i]$, or (ii) $R_j(t)\bar{u}_j(t) = 0$ for all $t \in [0, \tau_i]$.

(i) Here, either $\bar{\beta}_{ji} = 0$; or $(S_i(t) + \pi_{ji}I_i(t)) = 0$, and hence due to Theorem 1, $S_i(t) = 0$ and $\pi_{ji}I_i(t) = 0$. In the latter case, from Theorem 1, $S_i(0) = 0$ and $\pi_{ji} = 0$. and therefore for all $t > 0$, we will have $\bar{\beta}_{ji}(S_i + \pi_{ji}I_i)R_j\bar{u}_j = 0$.

(ii) We can assume $\bar{\beta}_{ji}(S_i(t) + \pi_{ji}I_i(t)) > 0$ for all $t \in (0, \tau_i]$. For such j, if $\tau_j < \tau_i$, $R_j(t) > 0$ for $t \in (\tau_j, \tau_i]$, therefore $\bar{u}_j(t) = 0$ for such t. Due to the structure results obtained for the interval $[\tau_j, T]$ in Theorem 3, $\bar{u}_j(t) = 0$ for all $t > \tau_j$, and therefore $\bar{\beta}_{ji}(S_i + \pi_{ji}I_i)R_j\bar{u}_j = 0$ for all $t > 0$.

Now, since $M < \infty$, the set $W = \{\tau_k : \tau_k > 0, k = 1, \ldots, M\}$ must have a minimum $\omega_0 < T$. Let $L(\omega_0) = \{k \in \{1, \ldots, M\} : \tau_k = \omega_0\}$. Let the second smallest element in W be ω_1. Using the above argument, the values of $R_k(t)$ for $t \in [\omega_0, \omega_1]$ for all $k \in L(\omega_0)$ would affect each other, but not R_i's for i such that $\tau_i > 0$, $i \notin L(\omega_0)$. Furthermore, in this interval for $k \in L(\omega_0)$ we have $\dot{R}_k = \sum_{g\in L(\omega_0)}\bar{\beta}_{gk}(S_g + \pi_{gk}I_i)R_g\bar{u}_g$, with $R_k(\omega_0) = 0$. We see that for all $k \in L(\omega_0)$, replacing $R_k(t) = 0$ in the RHS of equation (5) gives us $\dot{R}_k(t) = 0$, a compatible LHS, while not compromising the existence of solutions for all other states. An application of Theorem 1 for $t \in [\omega_0, \omega_1]$ and $\bar{u}$ shows that this is the unique solution of the system of differential equations (5). This contradicts the definition of τ_k, completing the proof of the lemma. ■

V. An Alternative Cost Functional

Recall that in our objective function, the cost of non-replicative patching was defined as $\sum_{i=1}^{M} R^0_i h_i(u_i)$ (respectively $\sum_{i=1}^{M} R_i h_i(u_i)$ for the replicative case), which corresponds to a scenario in which the dispatchers are charged for every instant they are immunizing/healing (distributing the patch), irrespective of the number of nodes they are delivering patches to. This represents a *broadcast* cost model where each transmission can reach all nodes of the neighbouring types. In an alternative *unicast* scenario, different transmissions may be required to deliver the patches to different nodes. This model is particularly useful if the dispatchers may only transmit to the nodes that have not yet received the patch.[10] Hence, the cost of patching in this case can be represented by: $\sum_{i=1}^{M}\sum_{j=1}^{M} R^0_i\bar{\beta}_{ij}(S_j + I_j)p(u_i)$ (for the replicative case: $\sum_{i=1}^{M}\sum_{j=1}^{M} R_i\bar{\beta}_{ij}(S_j + I_j)p(u_i)$), where $p(.)$ is an increasing function. More generally, the patching cost can be represented as a sum of the previously seen cost (§II-E) and this term.

[9]The RHS of the equation is evaluated at $t = t^*$ due to continuity.

[10]This can be achieved by keeping a common database of nodes that have successfully received the patch, or by implementing a turn-taking algorithm preventing double targeting. Note that we naturally assume that the network does not know with a priori certainty which nodes are infective, and hence it cannot differentiate between susceptibles and infectives. Consequently, even when $\pi_{ij} = 0$, i.e., the system manager knows the patch cannot remove the infection and can only immunize the susceptible, still the best it may be able to do is to forward the message to any node that has not yet received it.

For non-replicative patching, if all $h_i(\cdot)$ and $p(\cdot)$ are concave, then Theorem 2 will hold if for all pairs (i,j), $\pi_{ij}=\pi_j$ (i.e., healing efficacy only depends on the type of an infected node, not that of the immunizer). The analysis will change in the following ways: A term of $R_i^0 p(u_i)\sum_{j=1}^M \bar{\beta}_{ij}(S_j+I_j)$ is added to (15), and subsequently to (18) (with $u_i=1$ in the latter case). Also, (17) is modified by the subtraction of $\sum_{j=1}^M \bar{\beta}_{ji}R_j^0 p(u_j)$ from the RHS of both equations. This leaves $\dot{\lambda}_i^I-\dot{\lambda}_i^S$ untouched, while subtracting a positive amount from $\dot{\lambda}_i^I$, meaning that Lemma 1 still holds. As $\phi_i(t)$ was untouched, this means that Lemma 2 will also hold. Thus the RHS of $\dot{\gamma}_i$ is only modified by the subtraction of $\sum_{j,k=1}^M \bar{\beta}_{ij}(S_j+\pi_j I_j)\bar{\beta}_{kj}R_k^0\left(p(u_k)-u_k p(1)\right)$ which is a positive term, as for any continuous, increasing, concave function $p(\cdot)$ such that $p(0)=0$, we have $ap(b)\geq bp(a)$ if $a\geq b\geq 0$, since $\frac{p(x)}{x}$ is increasing. This yields: $(p(u_k)-u_k p(1)\geq 0)$. Therefore the conclusion of Theorem 2 holds. Similarly, it may be shown that Theorem 2 also holds for strictly convex $h_i(\cdot)$ provided $p(\cdot)$ is linear.

For the replicative case, if $p(\cdot)$ is linear ($p(x)=Cx$) and again $\pi_{ij}=\pi_j$ for all (i,j), Theorem 3 will hold. The modifications of the integrand and ψ_i are as above. The adjoint equations (30) are modified by the subtraction of $\sum_{j=1}^M C\bar{\beta}_{ji}R_j u_j$ from $\dot{\lambda}_i^I$ and $\dot{\lambda}_i^I$, and the subtraction of $Cu_i\sum_{j=1}^M \bar{\beta}_{ij}(S_j+I_j)$ from $\dot{\lambda}_i^R$. Due to the simultaneous change in ψ_i, however, we still have $\dot{\lambda}_i^R=\partial L(\mathbf{R})/\partial R_i-\alpha_i(u_i,t)$. Therefore, Lemma 3 still holds, as $\dot{\lambda}_i^I-\dot{\lambda}_i^S$ is unchanged, and a positive amount is subtracted from $\dot{\lambda}_i^I-\dot{\lambda}_i^R$. We absorb $\sum_{j=1}^M C\bar{\beta}_{ij}(S_j+I_j)$ into $\phi_i(t)$, where all the $p(\cdot)$ terms in $\dot{\phi}_i$ will cancel out, leaving the rest of the analysis, including for Lemmas 4 and 5, to be the same. The theorem follows.

VI. Numerical Investigations

In this section, we numerically investigate the optimal control policies for a range of malware and network parameters.[11] Recalling the notion of topologies presented in §II-A (in the paragraph before (2)), we consider three topologies: *linear*, *star* and *complete*, as was illustrated in Fig. 1. In our simulations, we assume that at $t=0$, only one of the regions (types) is infected, i.e., $I_i^0>0$ only for $i=1$. Also, $R_i^0=0.2$, $\beta_{ii}=\beta=0.223$ for all i.[12] The value of β_{ij}, $i\neq j$ is equal to $X_{Coef}\cdot\beta$ if link ij is part of the topology graph, and zero otherwise. (Unless otherwise stated, we use $X_{Coef}=0.1$.) It should be noted that $\beta_{ij}*T$ denotes the average number of contacts between nodes of regions i and j within the time period, and thus β and T are dependent variables. For simplicity, we use equal values for $\beta_{ji},\beta_{ij},\bar{\beta}_{ij},\bar{\beta}_{ji}$ for all i,j (i.e., $\beta_{ji}=\beta_{ij}=\bar{\beta}_{ij}=\bar{\beta}_{ji}$), and set $\pi_{ij}=\pi$ for all i,j. We examine two different aggregate cost structures for non-replicative patching:[13] (type-A)- $\int_0^T\left(K_I\sum_{i=1}^M I_i(t)+K_u\sum_{i=1}^M R_i^0 u_i(t)\right)dt$ and (type-B)- $\int_0^T\left(K_I\sum_{i=1}^M I_i(t)+K_u\sum_{i=1}^M R_i^0 u_i(t)(S_i(t)+I_i(t))\right)dt$ (described in §II-E and §V respectively). We select $T=35$, $K_I=1$, $K_u=0.5$ unless stated otherwise. For replicative patching, R_i^0 in both cost types is replaced with $R_i(t)$.

We first present an example of our optimal policy (§VI-1), and then we examine its behaviour in the linear and star topologies (§VI-2). Subsequently, we show the cost improvements it achieves over heuristics (§VI-3). Finally, we demonstrate the relative benefits of replicative patching (§VI-4).

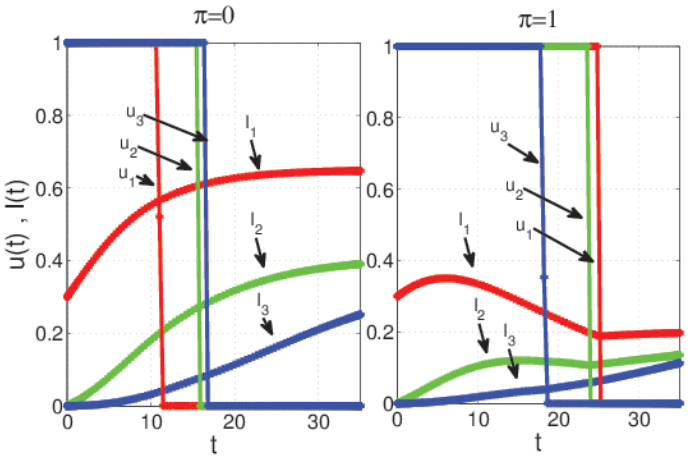

Fig. 2. Optimal patching policies and corresponding levels of infection in a three region linear topology. Note how the infection that initially only exists in region 1 spreads in region 1 and then to region 2, and finally to region 3.

[11] For our calculations, we use a combination of *C* programming and *PROPT®*, by *Tomlab Optimization Inc* for *MATLAB®*.

[12] This specific value of β is chosen to match the average inter-meeting times from the numerical experiment reported in [31].

[13] $f_i(\cdot)$, $h_i(\cdot)$, and $p(\cdot)$ are linear and identical for all i, and $l_i(\cdot)=0$.

1) Numeric Example: First, with the intention of illustrating our analytical results, in Fig. 2 we have depicted an example of the optimal dynamic patching policy along with the corresponding evolution of the infection as a function of time for a simple 3-region linear topology where the infection starts in region 1 ($\mathbf{I}^0 = (0.3, 0, 0)$). The cost model is type-A and patching is non-replicative. For $\pi = 0$ the levels of infection are non-decreasing, whereas for $\pi = 1$ they may go down as well as up (due to healing).

2) Effects of Topology: We study the *drop-off* times (the time thresholds at which the bang-bang optimal patching halts) in different regions for linear and star topologies.

Fig. 3 reveals two different patterns for $\pi = 0$ and $\pi = 1$ in a linear topology with 10 regions with non-replicative patching and type-A cost. For $\pi = 0$, a middle region is patched for the longest time, whereas for $\pi = 1$, as we move away from the origin of the infection (region 1), the drop-off point decreases. This is because for $\pi = 0$, patching can only benefit the network by recovering susceptibles. In regions closer to the origin, the fraction of susceptibles decreases quickly, making continuation of the patching comparatively less beneficial. In the middle regions, where there are more salvageable susceptibles, patching should be continued for longer. For regions far from the origin, patching can be stopped earlier, as the infection barely reaches them within the time horizon of consideration. For $\pi = 1$, patching is able to recover both susceptible and infective nodes. Hence, the drop-off times depend only on the exposure to the infection, which decreases with distance from the origin. As X_{Coef} is increased, the drop-off points when $\pi = 1$ get closer together. Intuitively, this is because higher cross-mixing rates have a homogenizing effect, as the levels of susceptible and infective nodes in different region rapidly become comparable. Also, Fig. 3 reveals that as X_{Coef} increases and more infection reaches farther regions, they are patched for longer, which agrees with our intuition.

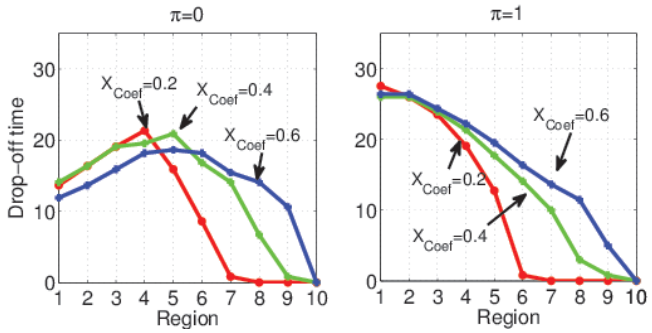

Fig. 3. Drop-off times in a linear topology for $X_{Coef} = 0.2, 0.4, 0.6$.

We next investigate a star configuration where the infection starts from a peripheral region (region 1), cost is type-B, patching is non-replicative, and $I_1^0 = 0.6$. Fig. 4 reveals the following interesting phenomenon: although the central region is the only one that is connected to all the regions, for $\pi = 0$, it is patched for shorter lengths of time compared to the peripherals. In retrospect, this is because only susceptible nodes can be patched and their number at the central region drops quickly due to its interactions with all the peripheral regions, rendering patching inefficient relatively swiftly. As expected, this effect is amplified with higher numbers of peripheral regions. For $\pi = 1$, on the other hand, the central region is patched for the longest time. This is because the infective nodes there can infect susceptible nodes in all regions, and hence the patching, which can now heal the infectives as well, does not stop until it heals almost all of infective nodes in this region.

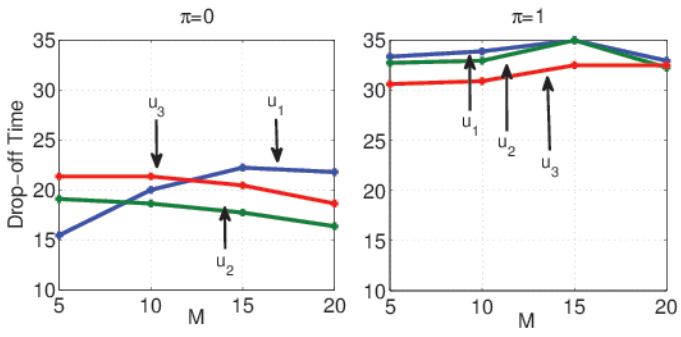

Fig. 4. Drop-off times in the star topology.

3) Cost Comparison: Next, in order to evaluate the efficacy of our dynamic heterogeneous patching policy, we compare our aggregate cost against those of four alternative patching policies. We label our policy as *Stratified Dynamic*. In the simplest alternative policy, all regions use identical patching intensities that do not change with time. We then select this fixed and static

level of patching so as to minimize the aggregate cost among all possible choices. We refer to this policy as *Static* (St.). The aggregate cost may be reduced if the static value of the patching is allowed to be distinct for different regions. These values (still fixed over time) are then independently varied and the best combination is selected. We refer to this policy as *Stratified Static* (S. St.). The third policy we implement is a *homogeneous* approximation to the heterogeneous network. Specifically, the whole network is approximated by a single region model with an equivalent inter-contact rate. This value is selected such that the *average* pairwise contact rates are equal in both systems. The optimal control is derived based on this model and applied across all regions to calculate the aggregate cost. We call this policy *Simplified Homogeneous* (S. H.). The simplified homogeneous policy is a special case of Spatially Static (Sp. St.) policies, where a one one-jump bang-bang control is applied to all regions to find the optimum uniform control.

Fig. 5 depicts the aggregate costs of all five policies for a linear topology with $M = 2, \ldots, 5$ regions. The cost is type-A and patching is replicative, with $\pi = 1$. Here, $I_1^0 = 0.2$, $K_u = 0.2$ and the rest of parameters are as before. As we can clearly observe, our stratified policy achieves the least cost, outperforming the rest. When the number of regions is small, H. and Sp. St. perform better than S. St., all of which obviously outperform St. However, as the number of regions increases and the network becomes more spatially heterogeneous, the homogeneous approximation, and all uniform controls in general, worsen and the S. St. policy quickly overtakes them as the best approximation to the optimal. For example for $M = 5$ regions, our policy outperforms the best static policies by 40% and the homogeneous approximation by 100%, which shows that our results about the structure of the optimal control can result in large cost improvements. For $\pi = 0$, a similar performance gap is observed. Specifically, as discussed, optimal drop-off times for this problem should vary based on the distance from the originating region, a factor that the Sp. St., H., and St. policies ignore.

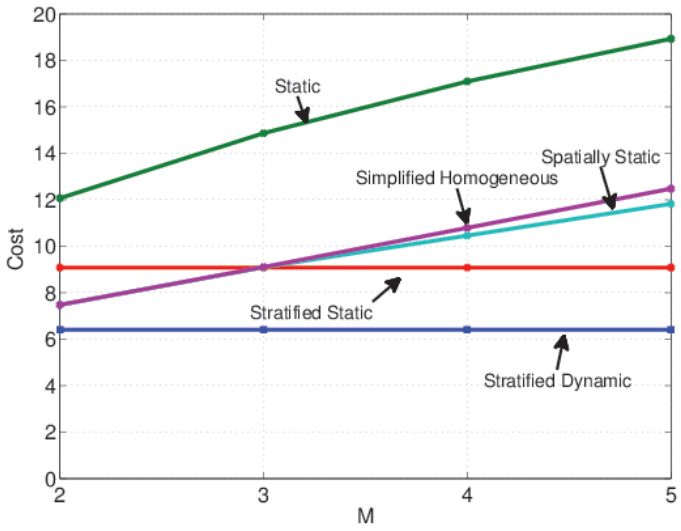

Fig. 5. Cost of heuristics vs. the optimal policy, linear topology.

4) Replicative vs. Non-replicative Patching: As previously stated, any solution to the non-replicative patching problem can be emulated by replicative patching, making this scenario worth investigation, even with the additional security vulnerabilities that the system has to contend with. In Fig. 6, we see the aggregate cost of optimal replicative and non-replicative patching in a complete topology as a function of the size of the network for $M \leq 11$, $\pi = 1$, and $K_u = 0.2$. Here, even for such modest sizes, replicative patching can be 60% more efficient than non-replicative patching, a significant improvement. This is especially true for the complete topology and other edge-dense topologies, as in replicative patching, the patch can spread in ways akin to the malware.

VII. Conclusion and Future Work

We considered the problem of disseminating security patches in a large resource-constrained heterogeneous network in the mean-field regime. Using tools from optimal control theory, we analytically proved that optimal dynamic policies for each type of node follow simple threshold-based structures. We numerically demonstrated the advantage of our heterogeneous policies over homogeneous approximations, as well as over static policies. For future research, we would like to further investigate the effects of heterogeneities in the structure of networks on both defense and attack strategies.

References

[1] K. Ramachandran and B. Sikdar, "On the stability of the malware free equilibrium in cell phones networks with spatial dynamics," in *ICC'07*, pp. 6169–6174, 2007.

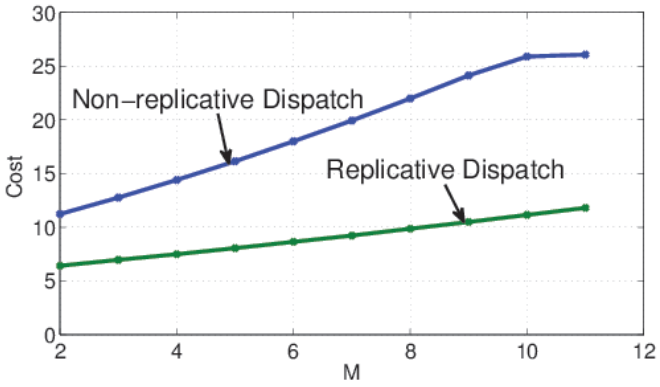

Fig. 6. Costs of replicative and non-replicative patching, complete topology.

[2] P. Wang, M. González, C. Hidalgo, and A. Barabási, "Understanding the spreading patterns of mobile phone viruses," *Science*, vol. 324, no. 5930, p. 1071, 2009.

[3] Z. Zhu, G. Cao, S. Zhu, S. Ranjan, and A. Nucci, "A social network based patching scheme for worm containment in cellular networks," in *INFOCOM'09*, pp. 1476–1484, IEEE, 2009.

[4] H. Leyden, "Blaster variant offers fix for pox-ridden pcs," *http://www.theregister.co.uk/2003/08/19/blaster_variant_offers_fix/*, 2003.

[5] M. Khouzani, S. Sarkar, and E. Altman, "Dispatch then stop: Optimal dissemination of security patches in mobile wireless networks," in *IEEE CDC'10*, pp. 2354–2359, 2010.

[6] M. Khouzani, S. Sarkar, and E. Altman, "Optimal control of epidemic evolution," in *IEEE INFOCOM*, 2011.

[7] J. Mickens and B. Noble, "Modeling epidemic spreading in mobile environments," in *Proceedings of the 4th ACM Workshop on Wireless Security*, pp. 77–86, ACM, 2005.

[8] K. Ramachandran and B. Sikdar, "Modeling malware propagation in networks of smart cell phones with spatial dynamics," in *IEEE INFOCOM'07*, pp. 2516–2520, 2007.

[9] Z. Chen and C. Ji, "Spatial-temporal modeling of malware propagation in networks," *IEEE Transactions on Neural Networks*, vol. 16, no. 5, pp. 1291–1303, 2005.

[10] F. Li, Y. Yang, and J. Wu, "CPMC: an efficient proximity malware coping scheme in smartphone-based mobile networks," in *IEEE INFOCOM'10*, pp. 1–9, 2010.

[11] Y. Yang, S. Zhu, and G. Cao, "Improving sensor network immunity under worm attacks: a software diversity approach," in *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pp. 149–158, ACM, 2008.

[12] Y. Li, P. Hui, L. Su, D. Jin, and L. Zeng, "An optimal distributed malware defense system for mobile networks with heterogeneous devices," *SECON 2011, IEEE*, pp. 1–9, 2011.

[13] H. Nguyen and Y. Shinoda, "A macro view of viral propagation and its persistence in heterogeneous wireless networks," in *Fifth International Conference on Networking and Services*, pp. 359–365, IEEE, 2009.

[14] M. Liljenstam, Y. Yuan, B. Premore, and D. Nicol, "A mixed abstraction level simulation model of large-scale internet worm infestations," in *10th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems, MASCOTS 2002*, pp. 109–116, IEEE, 2002.

[15] M. Faghani and H. Saidi, "Malware propagation in online social networks," in *4th IEEE International Conference on Malicious and Unwanted Software (MALWARE)*, pp. 8–14, 2009.

[16] G. Zyba, G. Voelker, M. Liljenstam, A. Méhes, and P. Johansson, "Defending mobile phones from proximity malware," in *INFOCOM 2009, IEEE*, pp. 1503–1511, IEEE, 2009.

[17] M. Altunay, S. Leyffer, J. Linderoth, and Z. Xie, "Optimal response to attacks on the open science grid," *Computer Networks*, 2010.

[18] G. Feichtinger, R. F. Hartl, and S. P. Sethi, "Dynamic optimal control models in advertising: recent developments," *Management Science*, vol. 40, no. 2, pp. 195–226, 1994.

[19] S. P. Sethi, "Dynamic optimal control models in advertising: a survey," *SIAM review*, vol. 19, no. 4, pp. 685–725, 1977.

[20] S. P. Sethi and G. L. Thompson, *Optimal control theory: applications to management science and economics*, vol. 101. Kluwer Academic Publishers Boston, 2000.

[21] H. Behncke, "Optimal control of deterministic epidemics," *Optimal control applications and methods*, vol. 21, no. 6, pp. 269–285, 2000.

[22] K. Wickwire, "Mathematical models for the control of pests and infectious diseases: a survey," *Theoretical Population Biology*, vol. 11, no. 2, pp. 182–238, 1977.

[23] J. Cuzick and R. Edwards, "Spatial clustering for inhomogeneous populations," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 73–104, 1990.

[24] W. Hsu and A. Helmy, "Capturing user friendship in WLAN traces," *IEEE INFOCOM poster*, 2006.

[25] T. Antunović, Y. Dekel, E. Mossel, and Y. Peres, "Competing first passage percolation on random regular graphs," *ArXiv e-prints*, 2011.

[26] T. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," *Journal of Applied Probability*, pp. 49–58, 1970.

[27] N. Gast, B. Gaujal, and J. Le Boudec, "Mean field for Markov decision processes: from discrete to continuous optimization," *Arxiv preprint arXiv:1004.2342*, 2010.

[28] B. Bollobás, *Modern graph theory*, vol. 184. Springer Verlag, 1998.

[29] R. F. Stengel, *Optimal control and estimation*. Dover, 1994.

[30] M. Khouzani, S. Sarkar, and E. Altman, "Maximum damage malware attack in mobile wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 20, no. 5, pp. 1347–1360, 2012.

[31] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *ACM SIGCOMM Workshop on Delay-tolerant Networking*, p. 251, ACM, 2005.

[32] S. Eshghi, M. Khouzani, S. Sarkar, N. B. Shroff, and S. S. Venkatesh, "Optimal energy-aware epidemic routing in dtns." Technical report, 2013. available at http://www.seas.upenn.edu/ swati/TACoptimalenergyrouting.pdf.

APPENDIX A
PROOF OF THEOREM 1

We use the following general result :

Lemma 6. *Suppose the vector-valued function* $\mathbf{f} = (f_i, 1 \le i \le 3M)$ *has component functions given by quadratic forms* $f_i(t,\mathbf{x}) = \mathbf{x}^T Q_i(t)\mathbf{x} + {p_i}^T\mathbf{x} \quad (t \in [0,T];\ \mathbf{x} \in \mathbb{S})$*, where* $\mathbb{S}$ *is the set of* $3M$*-dimensional vectors* $\mathbf{x} = (x_1, \ldots, x_{3M})$ *satisfying* $\mathbf{x} \ge \mathbf{0}$ *and* $\forall j \in \{1, \ldots, M\}; x_j + x_{M+j} + x_{2M+j} = 1$*,* $Q_i(t)$ *is a matrix whose components are uniformly, absolutely bounded over* $[0,T]$*, as are the elements of the vector* p_i*. Then, for an* $3M$*-dimensional vector-valued function* $\mathbf{F}$*, the system of differential equations*

$$\begin{gathered}\dot{\mathbf{F}} = \mathbf{f}(t,\mathbf{F}) \qquad (0 < t \le T) \\ \textit{subject to initial conditions } \mathbf{F}(0) \in \mathbb{S}\end{gathered} \tag{34}$$

has a unique solution, $\mathbf{F}(t)$*, which varies continuously with the initial conditions* $\mathbf{F}_0 \in \mathbb{S}$ *at each* $t \in [0,T]$*.*

Proof: This lemma is virtually identical to Lemma 1 of [32] for $N = 3M$, with the difference that here $f_i(t,\mathbf{x})$ has an additive ${p_i}^T\mathbf{x}$ term. Thus, we will only list the changes: As the Euclidean norm $\|Q_i(t)\mathbf{x}\|$ is uniformly bounded over $(t,\mathbf{x}) \in [0,T] \times \mathbb{S}$, there exists $C < \infty$ such that $\sup_{[0,T]\times\mathbb{S}} \|Q_i(t)\mathbf{x}\| \le C$. Also, $\|p_i\| \le H$ for some $H < \infty$ as all its elements are bounded. Now, for each t, we may write

$$f_i(t,\mathbf{x}) - f_i(t,\mathbf{y}) = \big(Q_i(t)\mathbf{x} + p_i\big)^T(\mathbf{x}-\mathbf{y}) + (\mathbf{x}-\mathbf{y})^T Q_i(t)\mathbf{y}$$

Taking absolute values of both sides, we obtain

$$|f_i(t,\mathbf{x}) - f_i(t,\mathbf{y})| \le \big|\big(Q_i(t)\mathbf{x}\big)^T(\mathbf{x}-\mathbf{y})\big| + \big|(\mathbf{x}-\mathbf{y})^T Q_i(t)\mathbf{y}\big| + \big|{p_i}^T\big(\mathbf{x}-\mathbf{y}\big)\big| \le (2C+H)\|\mathbf{x}-\mathbf{y}\| \ (t \in [0,T];\ \mathbf{x},\mathbf{y} \in \mathbb{S}),$$

by using the triangle and Cauchy-Schwarz inequalities. Hence

$$\|\mathbf{f}(t,\mathbf{x}) - \mathbf{f}(t,\mathbf{y})\| \le L\,\|\mathbf{x}-\mathbf{y}\| \qquad (t \in [0,T];\ \mathbf{x},\mathbf{y} \in \mathbb{S}),$$

and so $\mathbf{f}(t,\cdot)$ is Lipschitz over $\mathbb{S}$ where the Lipschitz constant $L = (2C+H)\sqrt{3M}$ may be chosen uniformly for $t \in [0,T]$.

The rest of the proof is exactly as in [32]. ∎

Proof of Theorem 1: We write $\mathbf{F}(0) = \mathbf{F}_0$, and in a slightly informal notation, $\mathbf{F} = \mathbf{F}(t) = \mathbf{F}(t,\mathbf{F}_0)$ to acknowledge the dependence of $\mathbf{F}$ on the initial value $\mathbf{F}_0$.

We first verify that $\mathbf{S}(t) + \mathbf{I}(t) + \mathbf{R}(t) = \mathbf{1}$ for all t in both cases. By summing the left and right sides of the system of equations (2) and the $\dot{R}_i$ equation that was left out (respectively the two sides of equations (5)), we see that in both cases for all i, $\big(\dot{S}_i(t) + \dot{I}_i(t) + \dot{R}_i(t)\big) = 0$, and, in view of the initial normalization $\big(S_i(0) + I_i(0) + R_i(0)\big) = 1$, we have $\big(S_i(t) + I_i(t) + R_i(t)\big) = 1$ for all t and all i.

We now verify the non-negativity condition. Let $\mathbf{F} = (F_1, \ldots, F_{3M})$ be the state vector in $3M$ dimensions whose elements are comprised of $(S_i, 1 \le i \le M)$, $(I_i, 1 \le i \le M)$ and $(R_i, 1 \le i \le M)$ in some order. The system of equations (2) can thus be represented as $\dot{\mathbf{F}} = \mathbf{f}(t,\mathbf{F})$, where for $t \in [0,T]$ and $\mathbf{x} \in \mathbb{S}$, the vector-valued function $\mathbf{f} = (f_i, 1 \le i \le 3M)$ has component functions $f_i(t,\mathbf{x}) = \mathbf{x}^T Q_i(t)\mathbf{x} + {p_i}^T\mathbf{x}$ in which (i) $Q_i(t)$ is a matrix whose non-zero elements are of the form $\pm\beta_{jk}$, (ii) the elements of $p_i(t)$ are of the form $\pm\bar{\beta}_{jk}R_j^0 u_j$ and $\pm\bar{\beta}_{jk}\pi_{jk}R_j^0 u_j$, whereas (5) can be represented in the same form but with (i) $Q_i(t)$ having elements $\pm\beta_{jk}$, $\pm\bar{\beta}_{jk}u_j$, and $\pm\bar{\beta}_{jk}\pi_{jk}u_j$, and (ii) $p_i = \mathbf{0}$. Thus, the components of $Q_i(t)$ are uniformly, absolutely bounded over $[0,T]$. Lemma 6 establishes that the solution $\mathbf{F}(t,\mathbf{F}_0)$ to the systems (2) and (5) is unique and varies continuously with the initial conditions $\mathbf{F}_0$; it clearly varies continuously with time. Next, using elementary calculus, we show in the next paragraph that if $\mathbf{F}_0 \in \mathbf{Int}\ \mathbb{S}$ (and, in particular, each component of $\mathbf{F}_0$ is positive), then each component of the solution $\mathbf{F}(t,\mathbf{F}_0)$ of (2) and (5) is positive at each $t \in [0,T]$. Since $\mathbf{F}(t,\mathbf{F}_0)$ varies continuously with $\mathbf{F}_0$, therefore $\mathbf{F}(t,\mathbf{F}_0) \ge \mathbf{0}$ for all $t \in [0,T]$, $\mathbf{F}_0 \in \mathbb{S}$, which completes the overall proof.

Accordingly, let the S_i, I_i, and R_i component of $\mathbf{F}_0$ be positive. Since the solution $\mathbf{F}(t,\mathbf{F}_0)$ varies continuously with time, there exists a time, say $t' > 0$, such that each component of $\mathbf{F}(t,\mathbf{F}_0)$ is positive in the interval $[0,t')$. The result follows trivially if $t' \ge T$. Suppose now that there exists $t'' < T$ such that each component of $\mathbf{F}(t,\mathbf{F}_0)$ is positive in the interval $[0,t'')$, and at least one such component is 0 at t''.

We first examine the non-replicative case. We show that such components can not be S_i for any i and subsequently rule out I_i and R_i for all i. Note that $u_j(t), I_j(t), S_j(t)$ are bounded in $[0,t'']$ (recall $(S_j(t) + I_j(t) + R_j(t)) = 1, S_j(t) \ge 0, I_j(t) \ge 0, R_j(t) \ge 0$ for all $j \in \{1,\ldots,M\}, t \in [0,t'']$). From (2a) $S_i(t'') = S_i(0)e^{-\int_0^{t''}\sum_{j=1}^M(\beta_{ji}I_j(t) + \bar{\beta}_{ji}R_j^0 u_j(t))\,dt}$. Since all $u_j(t), I_j(t)$ are bounded in $[0,t'']$, $S_i(0) > 0, R_j^0 \ge 0$, and $\beta_{ji}, \bar{\beta}_{ji} \ge 0$, therefore $S_i(t'') > 0$. Since $S_i(t) > 0$, $I_i(t) \ge 0$ for all $i, t \in [0,t'']$, and $\beta_{ji} \ge 0$, from (2b), $\dot{I}_i \ge -I_i\sum_{j=1}^M \pi_{ji}\bar{\beta}_{ji}R_j^0 u_j$ for all i in the interval $[0,t'']$. Thus, $I_i(t'') \ge I_i(0)e^{-\int_0^{t''}\sum_{j=1}^M \pi_{ji}(\bar{\beta}_{ji}R_j^0 u_j(t))\,dt}$. Since all $u_j(t), I_j(t), S_j(t)$ are bounded in $[0,t'']$, and $I_i(0) > 0$, $\bar{\beta}_{ji}, \pi_{ji} \ge 0$, it follows that $I_i(t'') > 0$ for all $i \ge 0$. Finally, $R_i(t'') > 0$ because $R_i(0) > 0$ and $\dot{R}_i(t) \ge 0$ from the above, so $R_i(t) \ge R_i^0$,

and $S_i(t) + I_i(t) \leq 1 - R_i^0$ for all t and i. This contradicts the definition of t'' and in turn implies that $\mathbf{F}(t, \mathbf{F}_0) > 0$ for all $t \in [0, T]$, $\mathbf{F}_0 \in \mathbf{Int}\ \mathbb{S}$.

The proof for the replicative case is similar, with the difference that R_i^0 is replaced with R_i, which is itself bounded.

Since the control and the unique state solution $\mathbf{S}(t)$, $\mathbf{I}(t)$ are non-negative, (2a, 5a) imply that $\mathbf{S}(t)$ is a non-increasing function of time. Thus, $S_j(t) = 0$ if $S_j(0) = 0$ for any j. Using the argument in the above paragraph and starting from a $t' \in [0, T)$ where $S_j(t') > 0$, $I_j(t') > 0$, or $R_j(t') > 0$, it may be shown respectively that $S_j(t) > 0$, $I_j(t) > 0$, and $R_j(t) > 0$ for all $t > t'$. All that remains to show now is:

Lemma 7. *There exists $\epsilon > 0$ such that $\mathbf{I}(t) > 0$ for $t \in (0, \epsilon)$.*

Let $d(i, j)$ be the *distance* from type j to type i (i.e., for all i, $d(i, i) = 0$ and for all pairs (i, j), $d(i, j) = 1+$ minimum number of types in a path from type j to type i). Now, define $d(i, U) := \min_{j \in U} d(i, j)$, where $U := \{\, i : I_i^0 > 0 \,\}$. Since we assumed that every type i is either in U or is connected to a type in U, $d(i, U) < M$ for all types i.

Let $\delta > 0$ be a time such that for all types i such that $d(i, U) = 0$ (the initially infected types), we have $I_i(t) > 0$ for $t \in [0, \delta)$. Thus, proving Lemma 8 below will be equivalent to proving Lemma 7, given an appropriate scaling of δ.

Lemma 8. *For all i and for all integers $r \geq 0$, if $d(i, U) \leq r$, then $I_i(t) > 0$ for $t \in (\frac{r}{M}\delta, T)$.*

Proof: By induction on r.

Base case: $r = 0$. If $d(i, U) = 0$, this means that the type is initially infected, and thus $I_i(t) > 0$ for $t \in (0, T)$ by definition. Therefore the base case holds.

Induction step: Assume that the statement holds for $r = 0, \ldots, k$ and consider $r = k + 1$. Since $(\frac{k+1}{M}\delta, T) \subset (\frac{k}{M}\delta, T)$, we need to examine types i such that $d(i, U) = k + 1$. In equation (2b) at $t = \frac{k+1}{M}\delta$, the first sum on the right involves terms like $I_j(\frac{k+1}{M}\delta)S_i(\frac{k+1}{M}\delta)$ where j is a neighbor of i, while the second sum involves terms like $I_i(\frac{k+1}{M}\delta)u_j(\frac{k+1}{M}\delta)$. Since $d(i, U) = k+1$, there exist neighbours j of i such that $d(j, U) = k$, and therefore $I_j(t) > 0$ for $t \in [\frac{k+1}{M}\delta, T)$ (by the induction hypothesis). Hence since $S_i^0 > 0$ and $\beta_{ji} > 0$ (i and j being neighbours), for such t, $\dot{I}_i(t) > -I_i(t) \sum_{j=1}^{M} \pi_{ji}\bar{\beta}_{ji}R_j^0 u_j(t) \geq -GI_i(t)$, where $G \geq 0$ is an upperbound on the sum (continuous functions are bounded on a closed and bounded interval). Thus $I_i(t) > I_i(\frac{k+1}{M}\delta)e^{-Kt} > 0$, completing the proof for $r = k + 1$. ■